

We demonstrate a framework for improving the availability of cluster based Internet services. Our approach models Internet services as a collection of interconnected

components, each possessing well defined interfaces and failure semantics. Such a decomposition allows designers to engineer high availability based on an understanding of the interconnections and isolated fault behavior of each component, as opposed to ad-hoc methods. In this work, we focus on using the entire commodity workstation ...

³ Pursuing failure: the distribution of program failures in a profile space

William Dickinson, David Leon, Andy Podgurski

September 2001 ACM SIGSOFT Software Engineering Notes , Proceedings of the 8th European software engineering conference held jointly with 9th ACM SIGSOFT international symposium on Foundations of software

engineering ESEC/FSE-9, Volume 26 Issue 5

Publisher: ACM Press

Full text available: pdf(304.58 KB)

Additional Information: full citation, abstract, references, citings, index terms

Observation-based testing calls for analyzing profiles of executions induced by potential test cases, in order to select a subset of executions to be checked for conformance to requirements. A family of techniques for selecting such a subset is evaluated experimentally. These techniques employ automatic cluster analysis to partition executions, and they use various sampling techniques to select executions from clusters. The experimental results support the hypothesis that with appropriate profil ...

Keywords: adaptive sampling, cluster analysis, cluster filtering, failure-pursuit sampling, multivariate data analysis, observation-based testing, software testing

4 <u>Technical papers: consistency management and quality assurance: Automated</u> support for classifying software failure reports

Andy Podgurski, David Leon, Patrick Francis, Wes Masri, Melinda Minch, Jiayang Sun, Bin

May 2003 Proceedings of the 25th International Conference on Software Engineering

Publisher: IEEE Computer Society

Full text available: pdf(1.06 MB) Additional Information: full citation, abstract, references, citings, index Publisher Site

This paper proposes automated support for classifying reported software failures in order to facilitate prioritizing them and diagnosing their causes. A classification strategy is presented that involves the use of supervised and unsupervised pattern classification and multivariate visualization. These techniques are applied to profiles of failed executions in order to group together failures with the same or similar causes. The resulting classification is then used to assess the frequency and s ...

5 Analysis and implementation of software rejuvenation in cluster systems

Kalyanaraman Vaidyanathan, Richard E. Harper, Steven W. Hunter, Kishor S. Trivedi June 2001 ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 2001 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '01, Volume 29 Issue 1

Publisher: ACM Press

Full text available: pdf(983.05 KB) Additional Information: full citation, abstract, references, citings

Several recent studies have reported the phenomenon of "software aging", one in which the state of a software system degrades with time. This may eventually lead to performance degradation of the software or crash/hang failure or both. "Software rejuvenation" is a pro-active technique aimed to prevent unexpected or unplanned outages due to aging. The basic idea is to stop the running software, clean its internal state and restart it. In this paper, we discuss software rejuvenation as applied to ...

6 Manageability, availability, and performance in porcupine: a highly scalable, cluster-

based mail service

Yasushi Saito, Brian N. Bershad, Henry M. Levy

August 2000 ACM Transactions on Computer Systems (TOCS), Volume 18 Issue 3

Publisher: ACM Press

Full text available: pdf(2.52 MB) Additional Information: full citation, abstract, references, index terms

This paper describes the motivation, design and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail

service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying ...

Keywords: cluster, distributed systems, email, group membership protocol, load balancing, replication

7 �	A failure and overload tolerance mechanism for continuous media servers Rajesh Krishnan, Dinesh Venkatesh, Thomas D. C. Little November 1997 Proceedings of the fifth ACM international conference on Multimedia Publisher: ACM Press Full text available: pdf(2.23 MB) Additional Information: full citation, references, index terms	
	Keywords : caching, clustered video servers, content insertion, fault tolerance, interactive video-on-demand, overload tolerance, rate adaptive stream merging, stream clustering	
8	Method for distributed transaction commit and recovery using Byzantine Agreement within clusters of processors C. Mohan, R. Strong, S. Finkelstein July 1985 ACM SIGOPS Operating Systems Review, Volume 19 Issue 3 Publisher: ACM Press Full text available: pdf(1.11 MB) Additional Information: full citation, abstract, references This paper describes an application of Byzantine Agreement [DoSt82a, DoSt82e, LyFF82] to distributed transaction commit. We replace the second phase of one of the commit algorithms of [MoLi83] with Byzantine Agreement, providing certain trade-offs and advantages at the time of commit and providing speed advantages at the time of recovery	
9	from failure. The present work differs from that presented in [DoSt82b] by increasing the scope (handling a general tree of processes, and multi-cluster transac Method for distributed transaction commit and recovery using Byzantine Agreement	
٥	C. Mohan, R. Strong, S. Finkelstein August 1983 Proceedings of the second annual ACM symposium on Principles of distributed computing Publisher: ACM Press	
	Full text available: pdf(939.80 KB) Additional Information: full citation, abstract, references, citings, index terms This paper describes an application of Byzantine Agreement [DoSt82a, DoSt82c, LyFF82]	

to distributed transaction commit. We replace the second phase of one of the commit algorithms of [MoLi83] with Byzantine Agreement, providing certain trade-offs and advantages at the time of commit and providing speed advantages at the time of recovery from failure. The present work differs from that presented in [DoSt82b] by increasing the scope (handling a general tree of processes, and multi-cluster tr ...

10 Partition testing, stratified sampling, and cluster analysis

Andy Podgurski, Charles Yang

December 1993 ACM SIGSOFT Software Engineering Notes, Proceedings of the 1st ACM SIGSOFT symposium on Foundations of software engineering SIGSOFT '93, Volume 18 Issue 5

Publisher: ACM Press

Full text available: pdf(1.35 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

We present a new approach to reducing the manual labor required to estimate software reliability. It combines the ideas of *partition testing* methods with those of *stratified sampling* to reduce the sample size necessary to estimate reliability with a given degree of precision. Program executions are stratified by using automatic *cluster analysis* to group those with *similar features*. We describe the conditions under which stratification is effective for estimating softw ...

11 <u>Fastpath Optimizations for Cluster Recovery in Shared-Disk Systems</u>

Randal Burns

November 2004 Proceedings of the 2004 ACM/IEEE conference on Supercomputing

Publisher: IEEE Computer Society

Full text available: pdf(176.70 KB) Additional Information: full citation, abstract

We describe the design and implementation of a clustering service for a high-performance, shared-disk file system. The service provides failure detection and recovery, reliableend-to-end messaging, and a centralized and recoverable management interface. We implement novel optimizations in the voting protocol that resolves cluster membership. Optimizations allow clusters to form as quickly as possible without introducing livelock or requiring timeout parameters to be tuned carefully. Our treatmen ...

12 Manageability, availability and performance in Porcupine: a highly scalable, cluster-

based mail service

Yasushi Saito, Brian N. Bershad, Henry M. Levy

December 1999 ACM SIGOPS Operating Systems Review, Proceedings of the seventeenth ACM symposium on Operating systems principles SOSP

Publisher: ACM Press

Full text available: pdf(1.62 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

This paper describes the motivation, design, and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying serv ...

13 A High Availability Clustering Solution

Phil Lewis

August 1999 Linux Journal

Publisher: Specialized Systems Consultants, Inc.

Full text available: html(34.77 KB) Additional Information: full citation, abstract, index terms

Mr. Lewis tells us how he designed and implemented a simple high-availability solution for his company

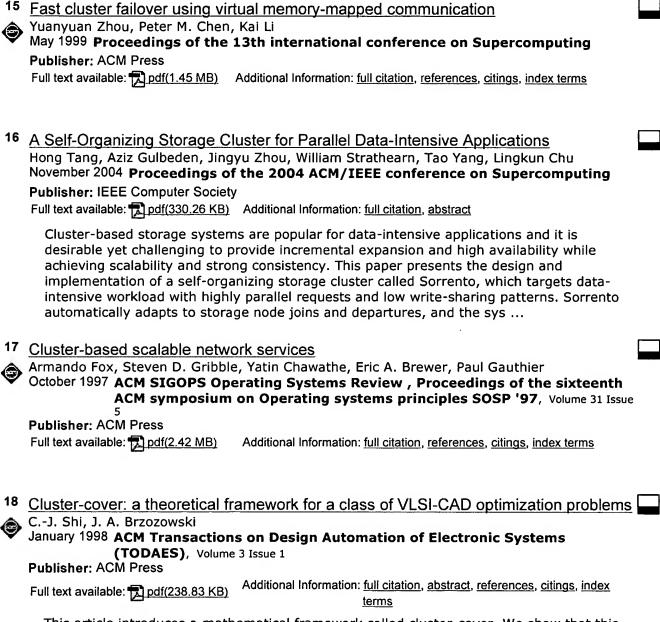
14 Quantifying and Improving the Availability of High-Performance Cluster-Based Internet Services

Kiran Nagaraja, Neeraj Krishnan, Ricardo Bianchini, Richard P. Martin, Thu D. Nguyen November 2003 **Proceedings of the 2003 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society

Full text available: pdf(306.01 KB) Additional Information: full citation, abstract

Cluster-based servers can substantially increase performance when nodes cooperate to globally manage resources. However, in this paper we show that cooperation results in a substantial availability loss, in the absence of high-availability mechanisms. Specifically, we show that a sophisticated cluster-based Web server, which gains a factor of 3 in performance through cooperation, increases service unavailability by a factor of 10 over a non-cooperative version. We then show how to augment this W ...



This article introduces a mathematical framework called cluster-cover. We show that this framework captures the combinatorial structure of a class of VLSI design optimization problems, including two-level logic minimization, constrained encoding, multilayer topological planar routing, application timing assignment for delay-fault testing, and minimization of monitoring logic for BIST enchancement. These apparently unrelated problems can all be cast into two metaproblems in our framework: fi ...

Keywords: NP-completeness, cluster-cover, logic minimization, self-checking logic design, state assignment, topological routing

19 Cellular disco: resource management using virtual clusters on shared-memory

multiprocessors

Kinshuk Govil, Dan Teodosiu, Yonggiang Huang, Mendel Rosenblum August 2000 ACM Transactions on Computer Systems (TOCS), Volume 18 Issue 3

Publisher: ACM Press

Full text available: pdf(287.05 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that laverages the existing operating system technology. In this paper we present a ...

Keywords: fault containment, resource managment, scalable multiprocessors, virtual machines

20 Cellular Disco: resource management using virtual clusters on shared-memory

<u>multiprocessors</u>

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

December 1999 ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles SOSP **'99**. Volume 33 Issue 5

Publisher: ACM Press

Full text available: pdf(1.93 MB)

Additional Information: full citation, abstract, references, citings, index

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a syste ...

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

The Guide

cluster "failure rate" availability "reboot rate"

SEARCH



Feedback Report a problem Satisfaction survey

Terms used cluster failure rate availability reboot rate

Found 3,834 of 167,655

Sort results

Save results to a Binder Search Tips

Try an Advanced Search Try this search in The ACM Guide

Display results

Best 200 shown

expanded form 🗢

Open results in a new window

Results 1 - 20 of 200

Result page: **1** 2 3 4 5 6 7 8 9 10

1 The costs and limits of availability for replicated services Relevance scale

Haifeng Yu, Amin Vahdat

October 2001 ACM SIGOPS Operating Systems Review, Proceedings of the eighteenth ACM symposium on Operating systems principles SOSP '01, Volume 35 Issue

Publisher: ACM Press

Full text available: pdf(1.46 MB)

Additional Information: full citation, abstract, references, citings, index terms

As raw system and network performance continues to improve at exponential rates, the utility of many services is increasingly limited by availability rather than performance. A key approach to improving availability involves replicating the service across multiple, wide-area sites. However, replication introduces well-known tradeoffs between service consistency and availability. Thus, this paper explores the benefits of dynamically trading consistency for availability using a continuous consi ...

Improving cluster availability using workstation validation



Taliver Heath, Richard P. Martin, Thu D. Nguyen

June 2002 ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '02, Volume 30 Issue 1

Publisher: ACM Press

Full text available: pdf(201.72 KB) Additional Information: full citation, abstract, references

We demonstrate a framework for improving the availability of cluster based Internet services. Our approach models Internet services as a collection of interconnected components, each possessing well defined interfaces and failure semantics. Such a decomposition allows designers to engineer high availability based on an understanding of the interconnections and isolated fault behavior of each component, as opposed to ad-hoc methods. In this work, we focus on using the entire commodity workstation ...

Analysis and implementation of software rejuvenation in cluster systems

Kalyanaraman Vaidyanathan, Richard E. Harper, Steven W. Hunter, Kishor S. Trivedi June 2001 ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 2001 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '01, Volume 29 Issue 1

Publisher: ACM Press

Full text available: pdf(983.05 KB) Additional Information: full citation, abstract, references, citings

Several recent studies have reported the phenomenon of "software aging", one in which the state of a software system degrades with time. This may eventually lead to performance degradation of the software or crash/hang failure or both. "Software rejuvenation" is a pro-active technique aimed to prevent unexpected or unplanned outages due to aging. The basic idea is to stop the running software, clean its internal state and restart it. In this paper, we discuss software rejuvenation as applied to ...

Survey of software tools for evaluating reliability, availability, and serviceability

Allen M. Johnson, Miroslaw Malek
September 1988 ACM Computing Surveys (CSUR), Volume 20 Issue 4

Publisher: ACM Press

Full text available: pdf(3.79 MB)

Additional Information: full citation, abstract, references, citings, index terms

In computer design, it is essential to know the effectiveness of different design options in improving performance and dependability. Various software tools have been created to evaluate these parameters, applying both analytic and simulation techniques, and this paper reviews those related primarily to reliability, availability, and serviceability. The purpose, type of models used, type of systems modeled, inputs, and outputs are given for each package. Examples of some of the key modeling ...

⁵ Multiview access protocols for large-scale replication

Xiangning Liu, Abdelsalam Helal, Weimin Du

June 1998 ACM Transactions on Database Systems (TODS), Volume 23 Issue 2

Publisher: ACM Press

Full text available: pdf(365.98 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

The article proposes a scalable protocol for replication management in large-scale replicated systems. The protocol organizes sites and data replicas into a tree-structured, hierarchical cluster architecture. The basic idea of the protocol is to accomplish the complex task of updating replicated data with a very large number of replicas by a set of related but independently committed transactions. Each transaction is responsible for updating replicas in exactly one cluster and invoking add ...

Keywords: data replication, large-scale systems, multiview access

6 Session 3: Minimal replication cost for availability

Haifeng Yu, Amin Vahdat

July 2002 Proceedings of the twenty-first annual symposium on Principles of distributed computing

Publisher: ACM Press

Full text available: pdf(1.18 MB) Additional Information: full citation, abstract, references, citings

Today, the utility of many replicated Internet services is limited by availability rather than raw performance. To better understand the effects of replica placement on availability, we propose the problem of minimal replication cost for availability. Let replication cost be the cost associated with replica deployment, dynamic replica creation and teardown at ncandidate locations. Given client access patterns, replica failure patterns, network partition patterns, a required consis ...

7 DNS: Availability, usage, and deployment characteristics of the domain name system



Jeffrey Pang, James Hendricks, Aditya Akella, Roberto De Prisco, Bruce Maggs, Srinivasan

October 2004 Proceedings of the 4th ACM SIGCOMM conference on Internet measurement

Publisher: ACM Press

Full text available: pdf(856.34 KB) Additional Information: full citation, abstract, references, index terms

The Domain Name System (DNS) is a critical part of the Internet's infrastructure, and is one of the few examples of a robust, highly-scalable, and operational distributed system. Although a few studies have been devoted to characterizing its properties, such as its workload and the stability of the top-level servers, many key components of DNS have not yet been examined. Based on large-scale measurements taken fromservers in a large content distribution network, we present a detailed study of ...

Keywords: DNS, availability, federated

8 A characterization of the simple failure-biasing method for simulations of highly

reliable Markovian Systems

Marvin K. Nakayama

January 1994 ACM Transactions on Modeling and Computer Simulation (TOMACS), Volume 4 Issue 1

Publisher: ACM Press

Full text available: pdf(2.25 MB)

Additional Information: full citation, abstract, references, citings, index terms

Simple failure biasing is an importance-sampling technique used to reduce the variance of estimates of performance measures and their gradients in simulations of highly reliable Markovian systems. Although simple failure biasing yields bounded relative error for the performance measure estimate when the system is balanced, it may not provide bounded relative error when the system is unbalanced. In this article, we provide a characterization of when the simple failure-biasing meth ...

Keywords: balanced failure biasing, gradient estimation, highly reliable systems, importance sampling, likelihood ratios, simple failure biasing

Bounding availability of repairable computer systems



R. R. Muntz, E. de Souza e Silva, A. Goyal

April 1989 ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 1989 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '89, Volume 17 Issue 1

Publisher: ACM Press

Full text available: pdf(1.15 MB)

Additional Information: full citation, abstract, references, citings, index terms

Markov models are widely used for the analysis of availability of computer/communication systems. Realistic models often involve state space cardinalities that are so large that it is impractical to generate the transition rate matrix let alone solve for availability measures. Various state space reduction methods have been developed, particularly for transient analysis. In this paper we present an approximation technique for determining steady state availability. Of particular interest is ...

10 Resource Management for Rapid Application Turnaround on Enterprise Desktop

Derrick Kondo, Andrew A. Chien, Henri Casanova

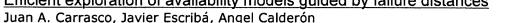
November 2004 Proceedings of the 2004 ACM/IEEE conference on Supercomputing

Publisher: IEEE Computer Society

Full text available: pdf(154.88 KB) Additional Information: full citation, abstract

Desktop grids are popular platforms for high throughput applications, but due their inherent resource volatility it is difficult to exploit them for applications that require rapid turnaround. Efficient desktop grid execution of short-lived applications is an attractive proposition and we claim that it is achievable via intelligent resource selection. We propose three general techniques for resource selection: resource prioritization, resource exclusion, and task duplication. We use these techni ...

11 Efficient exploration of availability models guided by failure distances



May 1996 ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 1996 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '96, Volume 24 Issue 1

Publisher: ACM Press

Full text available: pdf(1.08 MB)

Additional Information: full citation, abstract, references, index terms

Recently, a method to bound the steady-state availability using the failure distance concept has been proposed. In this paper we refine that method by introducing state space exploration techniques. In the methods proposed here, the state space is incrementally generated based on the contributions to the steady-state availability band of the states in the frontier of the currently generated state space. Several state space exploration algorithms are evaluated in terms of bounds quality and memor ...

12 Research papers: storage, indexing, and system architecture: Guaranteeing



Prakash Linga, Adina Crainiceanu, Johannes Gehrke, Jayavel Shanmugasudaram June 2005 Proceedings of the 2005 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(430.28 KB) Additional Information: full citation, abstract, references

New and emerging P2P applications require sophisticated range query capability and also have strict requirements on query correctness, system availability and item availability. While there has been recent work on developing new P2P range indices, none of these indices guarantee correctness and availability. In this paper, we develop new techniques that can provably guarantee the correctness and availability of P2P range indices. We develop our techniques in the context of a general P2P indexing ...

13 Cellular Disco: resource management using virtual clusters on shared-memory

multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

December 1999 ACM SIGOPS Operating Systems Review, Proceedings of the seventeenth ACM symposium on Operating systems principles SOSP

'99, Volume 33 Issue 5

Publisher: ACM Press

Full text available: pdf(1.93 MB)

Additional Information: full citation, abstract, references, citings, index terms

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology.In this paper we present a syste ...

14 Cellular disco: resource management using virtual clusters on shared-memory

multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3 Publisher: ACM Press

Full text available: pdf(287.05 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that laverages the existing operating system technology. In this paper we present a ...

Keywords: fault containment, resource managment, scalable multiprocessors, virtual machines

15 <u>Tunable randomization for load management in shared-disk clusters</u>

February 2005 ACM Transactions on Storage (TOS), Volume 1 Issue 1

Publisher: ACM Press

Full text available: pdf(551.85 KB) Additional Information: full citation, abstract, references, index terms

We develop and evaluate a system for load management in shared-disk file systems built on clusters of heterogeneous computers. It balances workload by moving file sets among cluster server nodes. It responds to changing server resources that arise from failure and recovery, and dynamically adding or removing servers. It also realizes performance consistency---nearly uniform performance across all servers. The system is adaptive and self-tuning. It operates without any *a priori* knowledge o ...

Keywords: Load management, computer clusters, heterogeneity, shared-disk file systems

¹⁶ I_{DDX}-based test methods: A survey

Sagar S. Sabade, Duncan M. Walker

April 2004 ACM Transactions on Design Automation of Electronic Systems (TODAES), Volume 9 Issue 2

Publisher: ACM Press

Full text available: pdf(1.83 MB) Additional Information: full citation, abstract, references, index terms

Supply current measurement-based test is a valuable defect-based test method for semiconductor chips. Both static leakage current (I_{DDQ}) and transient current (I_{DDT}) based tests have the capability of detecting unique defects that improve the fault detection capacity of a test suite. Collectively these test methods are known as I_{DDX} tests. However, due to advances in the semiconductor manufacturing process, the future of these test methods is uncertain. This pape ...

Keywords: I_{DDQ}, I_{DDT} test, VLSI testing, test

17 Understanding fault-tolerant distributed systems

Flavin Cristian

February 1991 Communications of the ACM, Volume 34 Issue 2

Publisher: ACM Press

Full text available: pdf(6.17 MB)

Additional Information: full citation, references, citings, index terms, review

18	Parallel architectures for processing high speed network signaling protocols
	Dipak Ghosal, T. V. Lakshman, Yennun Huang

December 1995 IEEE/ACM Transactions on Networking (TON), Volume 3 Issue 6

Publisher: IEEE Press

Full text available: pdf(1.58 MB)

Additional Information: full citation, references, citings, index terms

19 Methods for information server selection

David Hawking, Paul Thistlewaite

January 1999 ACM Transactions on Information Systems (TOIS), Volume 17 Issue 1

Publisher: ACM Press

Full text available: pdf(283.76 KB)

Additional Information: full citation, abstract, references, citings, index

The problem of using a broker to select a subset of available information servers in order to achieve a good trade-off between document retrieval effectiveness and cost is addressed. Server selection methods which are capable of operating in the absence of global information, and where servers have no knowledge of brokers, are investigated. A novel method using Lightweight Probe queries (LWP method) is compared with several methods based on data from past query processing, while Random and ...

Keywords: Lightweight Probe queries, information servers, network servers, server ranking, server selection, text retrieval

²⁰ Measurement and modeling of computer reliability as affected by system activity

R. K. Iyer, D. J. Rossetti, M. C. Hsueh

August 1986 ACM Transactions on Computer Systems (TOCS), Volume 4 Issue 3

Publisher: ACM Press

Full text available: pdf(1.44 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms, review

This paper demonstrates a practical approach to the study of the failure behavior of computer systems. Particular attention is devoted to the analysis of permanent failures. A number of important techniques, which may have general applicability in both failure and workload analysis, are brought together in this presentation. These include: smeared averaging of the workload data, clustering of like failures, and joint analysis of workload and failures. Approximately 17 percent of all failure ...

Results 1 - 20 of 200 Result page: 1 2 3 4 5 6 7 8 9 10

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Home | Login | Logout | Access Information | #

Welcome United States Patent and Trademark Office

Kampert K. -H.

■■Author Search

BROWSE

SEARCH

IEEE XPLORE GUIDE



OPTION 1

Quick Find an Author:

Enter a name to locate articles written by that author.

kampe

Example: Enter Lockett S to obtain a list of authors with the last name Lockett and the first initial S.

OPTION 2

Browse alphabetically

Select a letter from the list.

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

Select a name to view articles written by that author

Kampe A. Kampe G. Kampe M. Kampel M. Kampenes V. B. Kamper A. J. Kamper D. G. Kamper G. Kamper M. J. Kamper R. Kamperman A. Kamperman F. Kampers F. W. H. Kamperschroer J.

Kampes B.

Kampe J. Kampen G Kamper D. Kamper K. Kamper R. Kamperma Kampersct



Copyright 2005 I

Kampes B.

⊠e-mail



Home | Login | Logout | Access Information | Ale

Welcome United States Patent and Trademark Office

■ Search Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

>>

Results for "(kampe m.<in>au)"

Your search matched 2 of 1263585 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

View Session History

New Search

» Key

IEEE JNL

IEEE Journal or Magazine

IEE JNL

IEE Journal or Magazine

IEEE CNF **IEE CNF**

IEEE Conference Proceeding IEE Conference Proceeding

IEEE STD

IEEE Standard

Modify Search

(kampe m.<in>au)

Check to search only within this results set

Display Format:

Citation

Citation & Abstract

Select **Article Information**

1. The FAB predictor: using Fourier analysis to predict the outcome of conditional branches

Kampe, M.; Stenstrom, P.; Dubois, M.;

High-Performance Computer Architecture, 2002. Proceedings. Eighth International Symposium on

2-6 Feb. 2002 Page(s):223 - 232

AbstractPlus | Full Text: PDF(469 KB) IEEE CNF

2. Exploration of the spatial locality on emerging applications and the consequences for cache

Kampe, M.; Dahlgren, F.;

Parallel and Distributed Processing Symposium, 2000. IPDPS 2000. Proceedings. 14th Internation:

1-5 May 2000 Page(s):163 - 170

Digital Object Identifier 10.1109/IPDPS.2000.845978

AbstractPlus | Full Text: PDF(180 KB) | IEEE CNF

indexed by #Inspec Contact Us Privac

© Copyright 2005 IE